

BRIEF DESCRIPTION OF THE DRAWINGS:

[0012] The present invention is illustrated in the following drawings in which:

Fig. 1 is a schematic diagram of a network computing environment utilizing a channel subsystem usable with the present invention;

5 Fig. 2 is a schematic diagram of a single computer with shared physical memory and a plurality of discrete servers with a common lookup table of the present invention for transferring data from a sending-discrete server to a target-discrete server;

10 Fig. 3 is a schematic diagram illustrating the common lookup table of Fig. 2 including a hash tables control area, a source queue hash table, a target queue hash table, multiple queue controls, multiple QDIO queue sets, and means to add entries to the source queue hash table and target queue hash table;

Fig. 4 is a diagram of the hash tables control area of Fig. 3;

Fig. 5 is a diagram illustrating one of the queue controls of Fig. 3;

Fig. 6 is a diagram illustrating one of the queue sets of Fig. 3;

15 Fig. 7 is a diagram illustrating a send queue user buffer of the queue set of Fig. 6;

Fig. 8 is a diagram illustrating one of the entries of the source hash table of Fig. 3;

Fig. 9 is a diagram illustrating one of the entries of the target hash tables of Fig. 3;

Fig. 10 is an illustration of a three tier hierarchy of I/O completion vectors;

20 Fig. 11 is a schematic diagram of the hierarchy of Fig. 10 with a Time of Day (TOD) register, a Target Delay Interval (TDI) register, and a processor within a host computer for completion of I/O requests by devices;

Fig. 12 is a flow diagram showing the cooperation between the dispatcher of the OS and the devices;

25 Fig. 13 is a flow chart of the dispatcher program for an algorithm for determining the TDI value based upon workload heuristics;

Fig. 14 is a flow chart of the MakeDecision subroutine of the dispatcher program of Fig. 13;

Fig. 15 is a diagram illustrating a computer having multiple partitions, wherein interrupts of one of the partitions is under the control of a hypervisor; and

Fig. 16 is a diagram illustrating the use of an override bit for informing devices that immediate interrupts will be handled by the hypervisor of Fig. 15.

5 DESCRIPTION OF THE PREFERRED EMBODIMENT:

[0013] An example of an existing data processing system architecture is depicted in Fig. 1. As shown in Fig. 1, information is passed between the main storage 110, and one or more input/output devices (hereinafter I/O devices) 190, using channel subsystems 150. It will be understood that I/O devices as used herein refers to physical external I/O devices as well as virtual devices such as when data is transferred from one partition to another in an I/O manner and in which one partition appears as an I/O device to the other partition. In one embodiment, channel paths are established through the switch 160, the channel path comprising channels 155 and one or more control units shown at 180. These channel paths are the communication links established between the I/O devices 190 and the main storage for processing and exchange of information.

[0014] The main storage 110 stores data and programs which are input from I/O devices 190. Main storage is directly addressable and provides for high speed processing of data by central processing units and one or more I/O devices. One example of a main storage is a customer's storage area and a hardware system area (HSA) to be discussed later. I/O devices 190 pass information to or from main storage via facilities provided in the channel subsystem 250. Some examples of I/O devices include card readers and punches, magnetic-tape units, direct-access storage devices (DASD), displays, keyboards, printers, teleprocessing devices, communication controllers and sensor-based equipment.

[0015] The main storage is coupled to the storage control element (SCE) 120 which in turn is coupled to one or more central processing units (CPU) 130. The central processing unit(s) is the control center of the data processing system and typically comprises sequencing and processing facilities for instruction execution, initial program loading and other related functions. The CPU

is usually coupled to the SCE via a bi-directional or unidirectional bus. The SCE, which controls the execution and queuing of requests made by the CPU and channel subsystem, is coupled to the main storage, CPUs and the channel subsystem via different busses.

[0016] The channel subsystem directs the flow of information between I/O devices and main storage and relieves the CPUs of the task of communicating directly with the I/O devices so that data processing operations directed by the CPU can proceed concurrently with I/O processing operations. The channel subsystem uses one or more channel paths as the communication links in managing the flow of information to or from I/O devices. Each channel path consists of one or more channels, located within the channel subsystem, and one or more control units. In one preferred embodiment, a SAP I/O processor is also included as part of the channel subsystem.

[0017] As can be seen in Fig. 1, it is also possible to have one or more dynamic switches or even a switching fabric 195 (network of switches) included as part of the path, coupled to the channel(s) and the control unit(s). Each control unit is further attached via a bus to one or more I/O device(s).

[0018] The subchannel is the means by which the channel subsystem provides information about associated I/O devices to the central processing units; the CPUs obtain this information by executing I/O instructions. The subchannel consists of internal storage that contains information in the form of a channel command word (CCW) address, channel path identifier, device number, count, status indications, and I/O interruption subclass code, as well as information on path availability and functions pending or being performed. I/O operations are initiated with devices by executing I/O instructions that designate the subchannel associated with the device.

[0019] The execution of input/output operations is accomplished by the decoding and executing of CCWs by the channel subsystem and input/output devices. A chain of CCWs (input/output operations) is initiated when the channel transfers to the control unit the command specified by